

# **Opteron and AMD64**

## **A Commodity 64 bit x86 SOC**

Fred Weber

Vice President and CTO

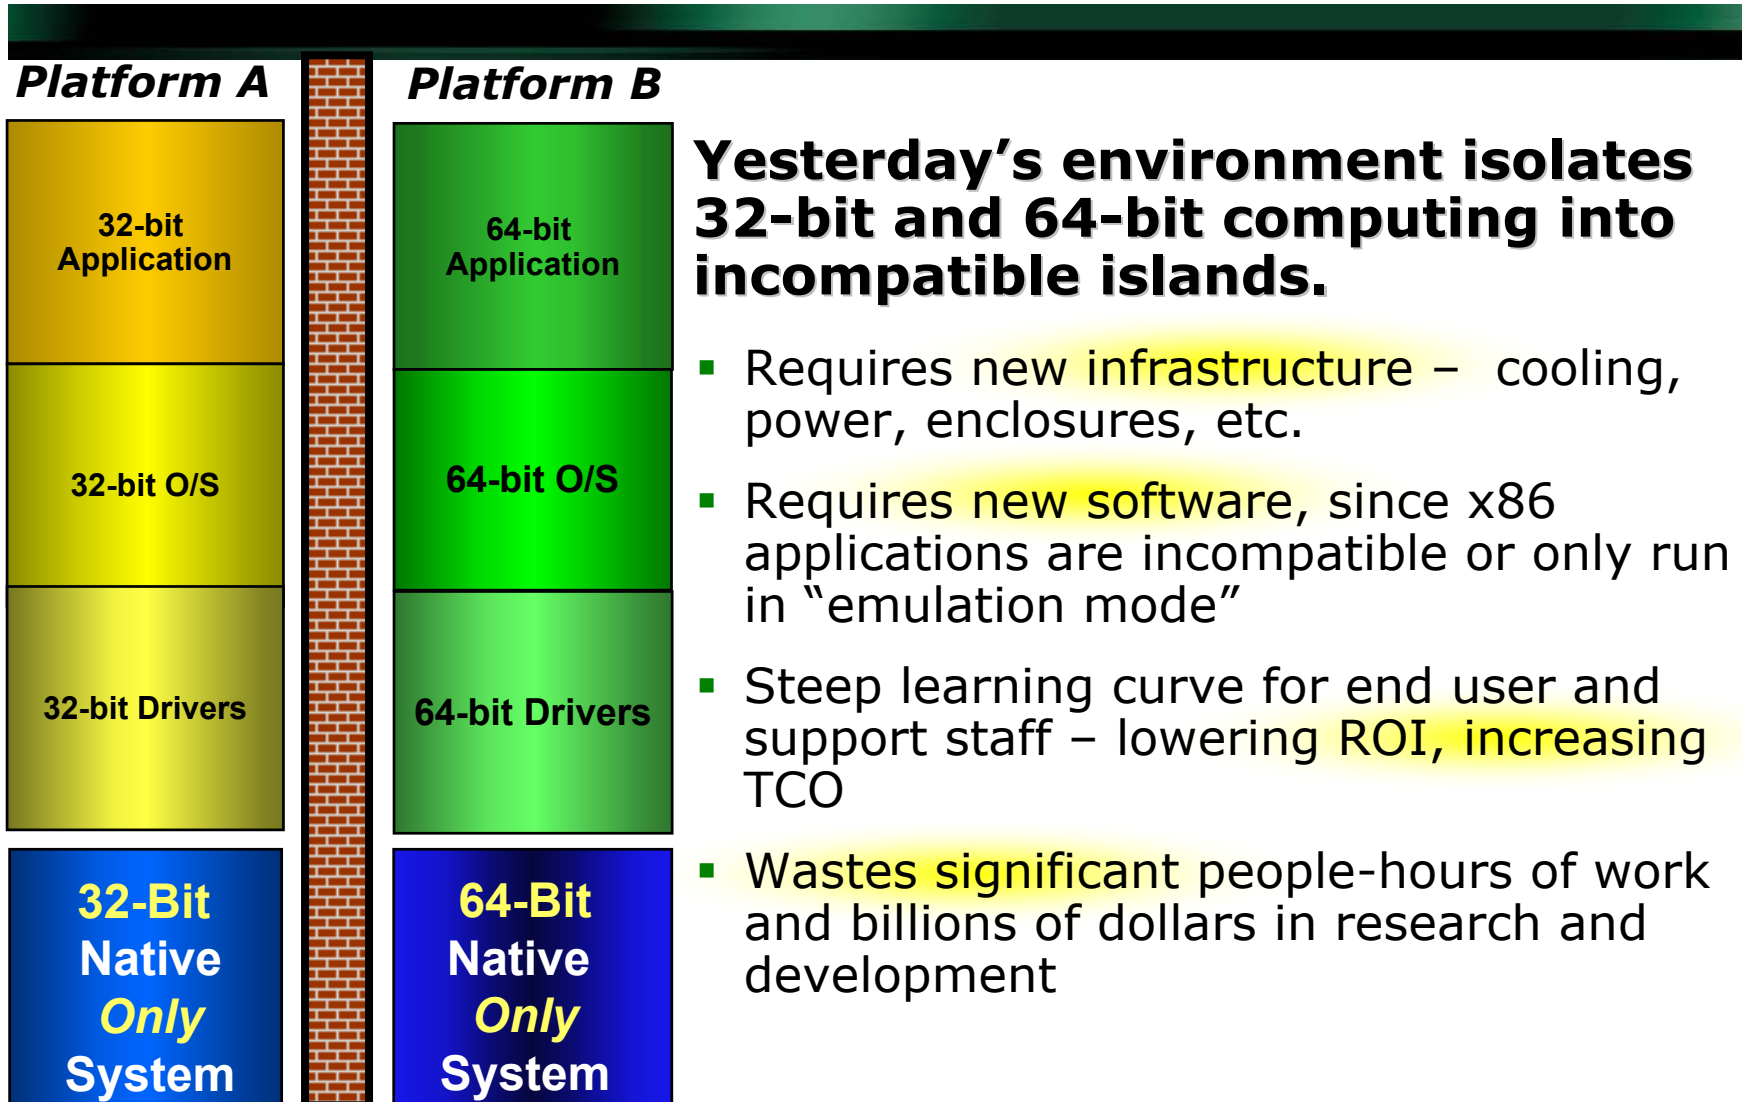
Computation Products Group  
Advanced Micro Devices

- Official Launch of AMD64 architecture and Production Server/Workstation CPUs
  - Series 200 (2P) available today
  - Series 800 (4P+) available later in Q2
- Oracle, IBM-DB2, Microsoft, RedHat, SuSe software support
  - And many others
- Dozens of server system vendors
  - System builder availability this quarter
  - IBM systems available 3Q03
- Lots of public benchmarks

# Before AMD64:



*Computing & infrastructure islands on either side of the wall*



# AMD's Industry Vision:

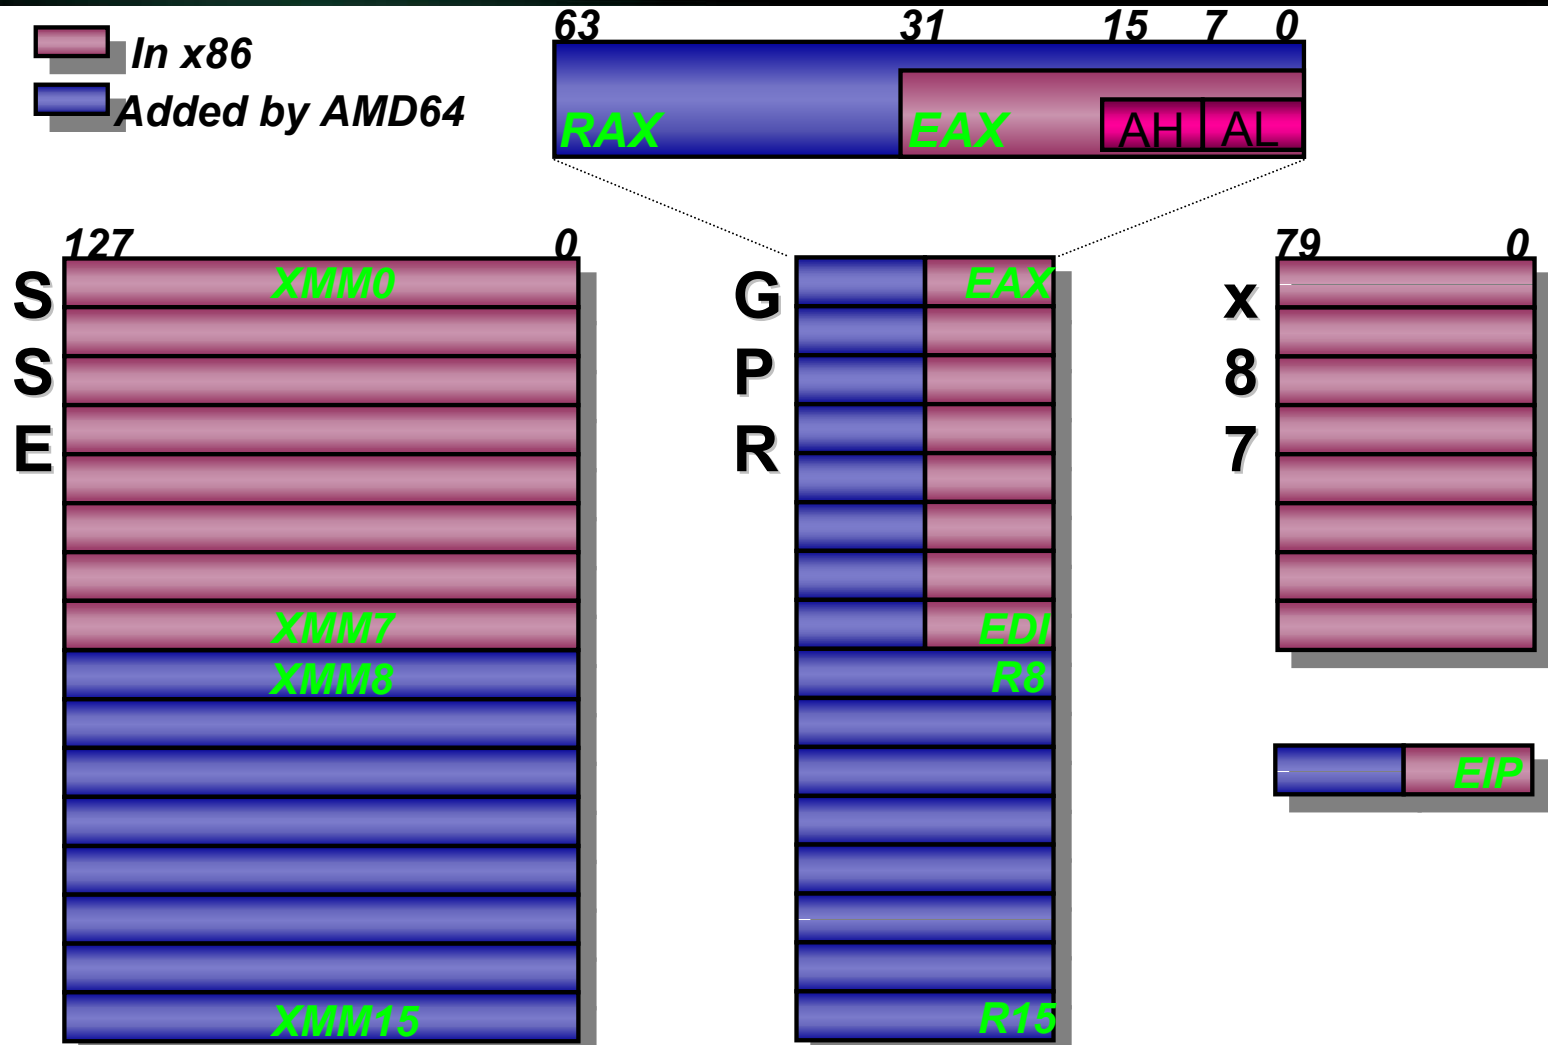
*Compatible systems that bridge from 32- to 64-bit*



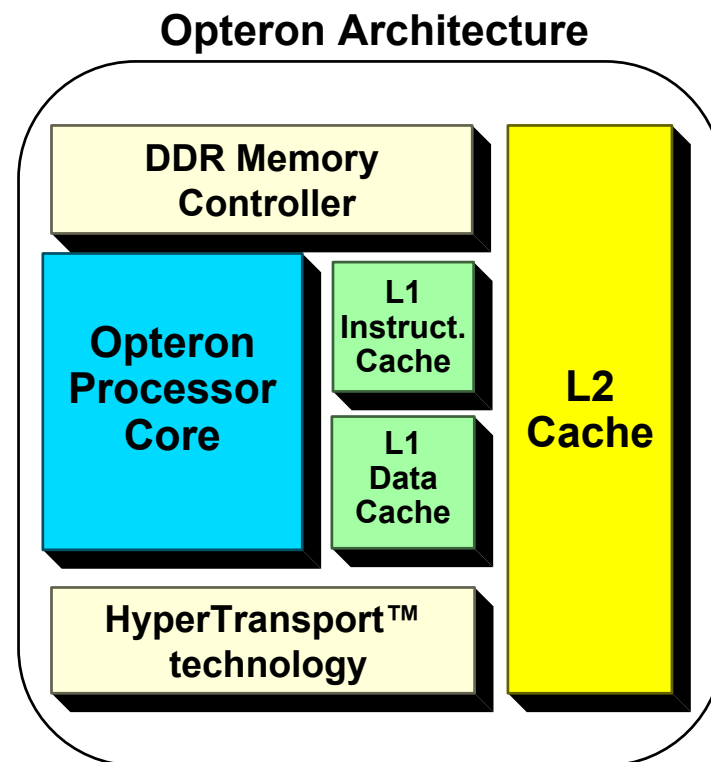
## AMD: Single Platform

|                                    |                                   |                         |
|------------------------------------|-----------------------------------|-------------------------|
| 32-Bit Application<br>(3 GB limit) | 32-Bit Application<br>(4GB limit) | 64-Bit Application      |
| 32-Bit O/S                         | 64-Bit O/S                        | 64-Bit O/S              |
| 32-Bit Software Drivers            | 64-Bit Software Drivers           | 64-Bit Software Drivers |
|                                    |                                   |                         |

- Leverages existing infrastructure
  - thermal, enclosures, power, and **BIOS**
- Runs existing 32-bit applications natively with unsurpassed performance
  - >20% increase clock-for-clock compared to AMP Athlon™ processor
  - No tools or O/S work needed
- Runs existing 32-bit applications on 64-bit O/S
  - Take full advantage of 4GB local memory
- Allows customers to migrate to 64-bit performance according to their schedule
- Low learning curve for users and support staff



- First AMD64 based processor
- Aggressive out-of-order, 9-issue superscalar processor
- Integrated DDR memory controller
- Leading performance in integer, floating point and multimedia
  - AMD64, x87, MMX™, 3DNow!™, SSE, SSE2
- Glueless multiprocessing through HyperTransport
- Expandable IO through HyperTransport



# AMD Opteron™ Processor Technology Overview



## • Processor Core Overview

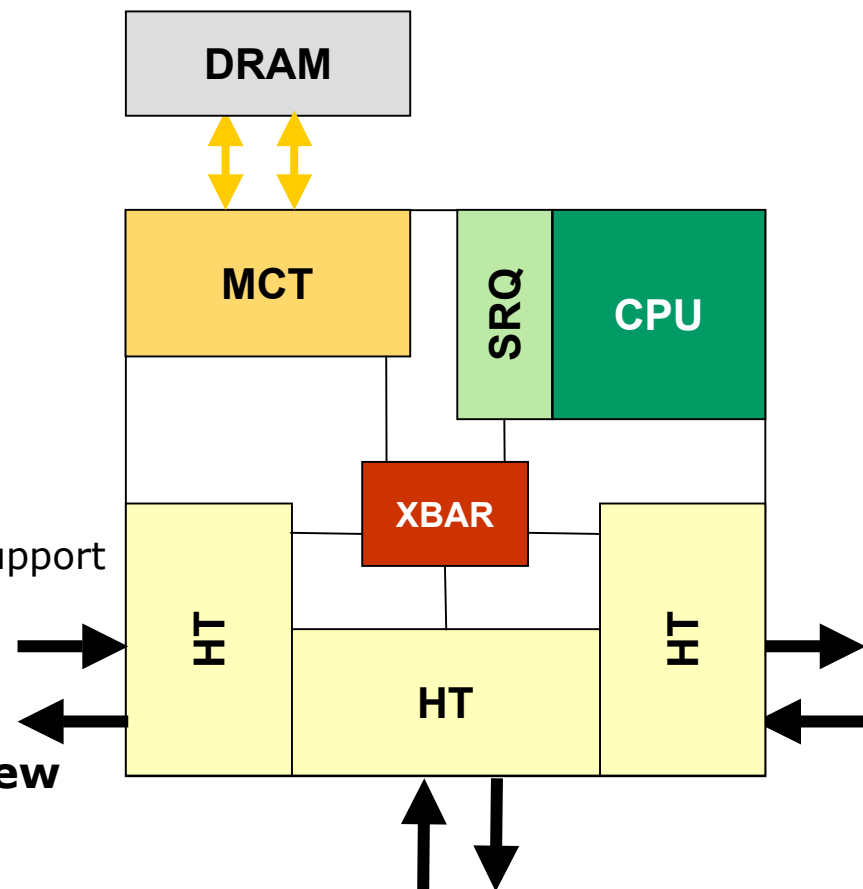
- Support for AMD's 64-bit technology
- 12-stage int, 17-stage fp pipelines
- Enhanced TLB structures
- TLB flush filter
- Enhanced branch prediction
- Large L2 cache (up to 1MB)
- ECC protection

## • Memory Controller Overview

- Dual-channel DDR memory
- PC2700, PC2100, or PC1600 DDR memory support
- Registered or Unbuffered DIMMs
- ECC and Chip Kill
- High bandwidth (up to **6.4GB/s**)

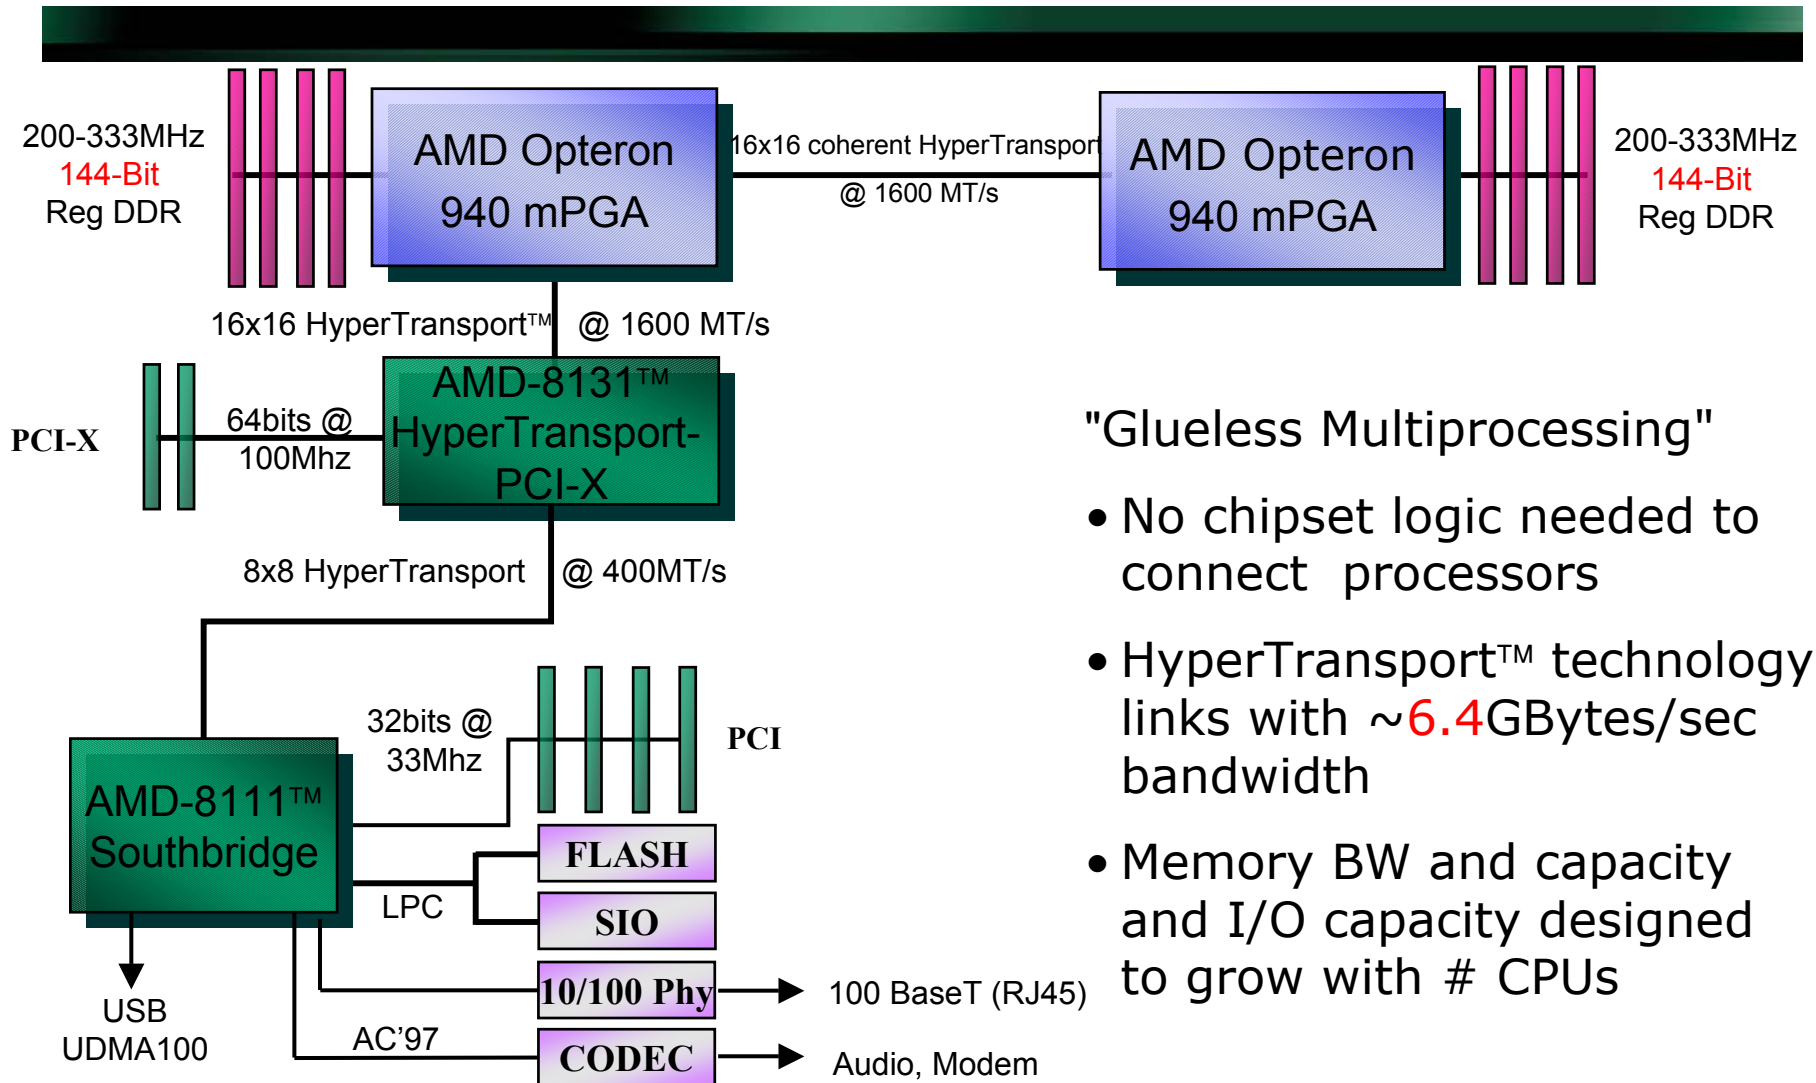
## • HyperTransport™ Technology Overview

- One, two, or three links
- 2, 4, 8, 16, or 32-bits full duplex
- Up to 6.4 GB/s bandwidth per link
- **19.2 GB/s** aggregate external bandwidth



HT = HyperTransport™ technology

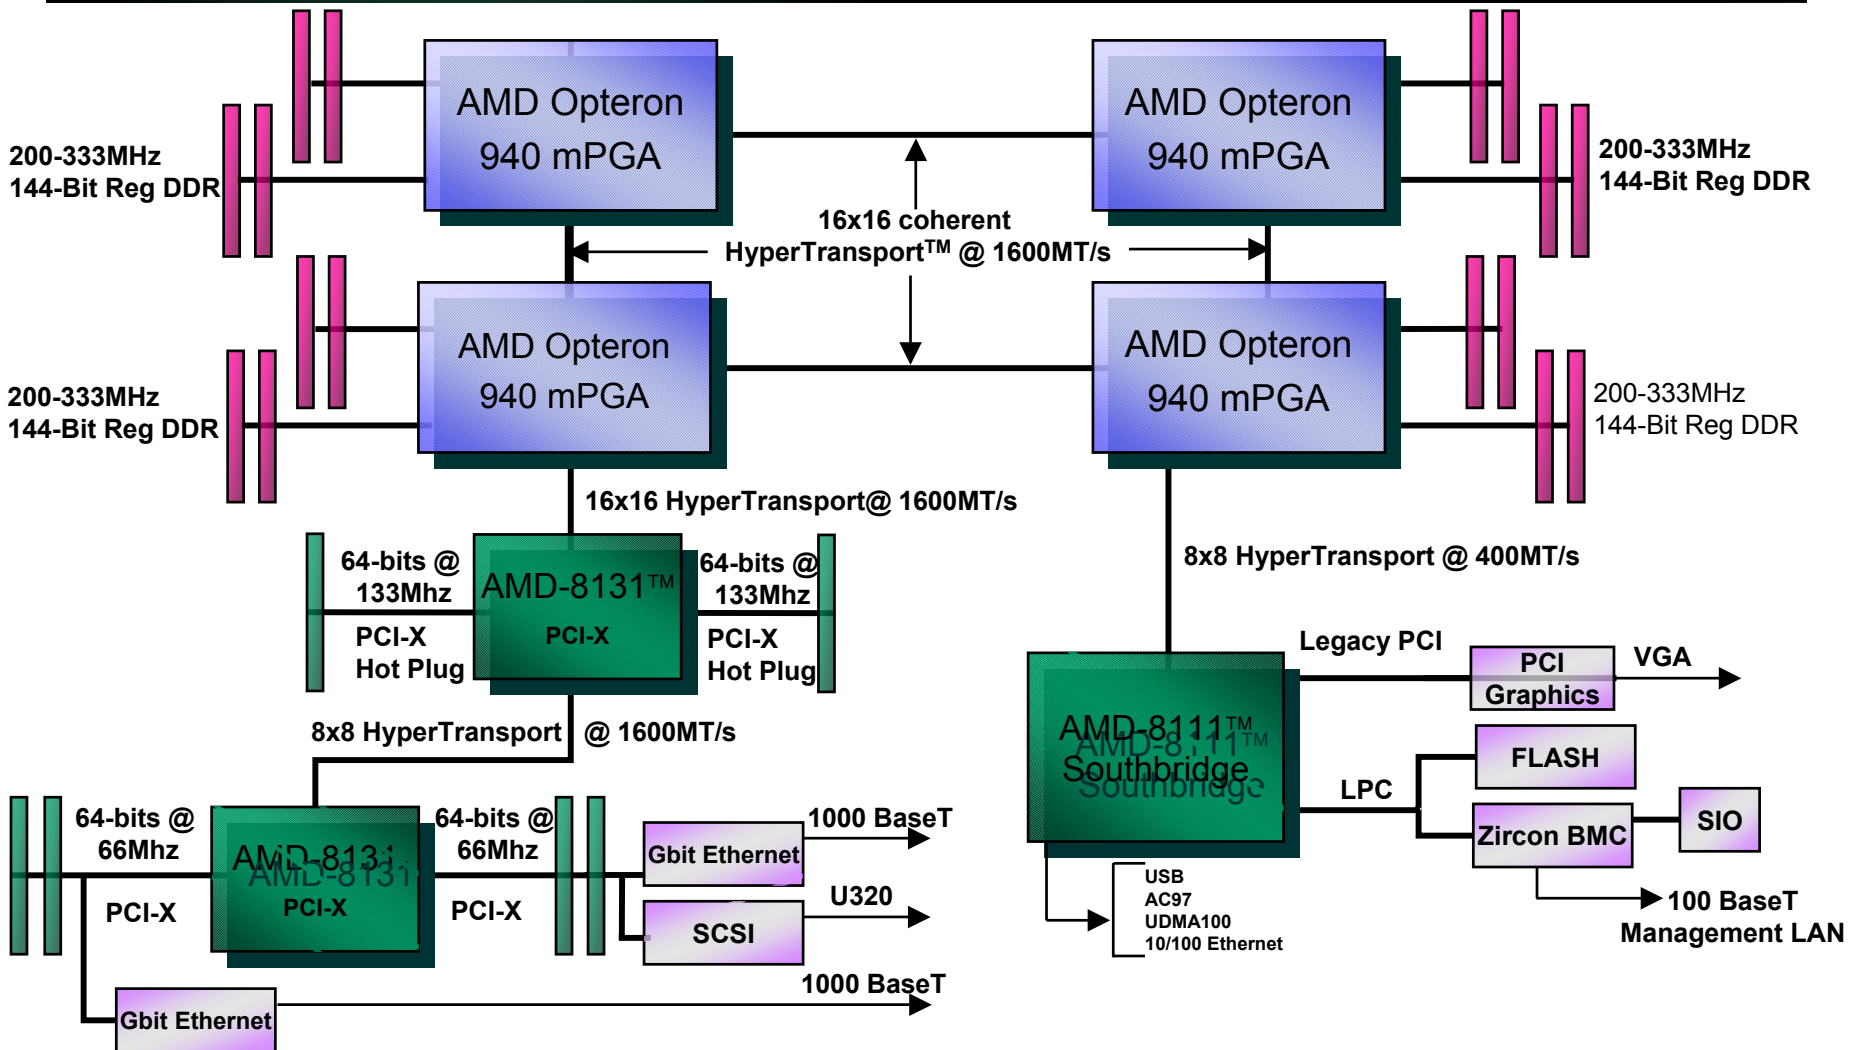
# AMD Opteron™ processor-based 2P Server



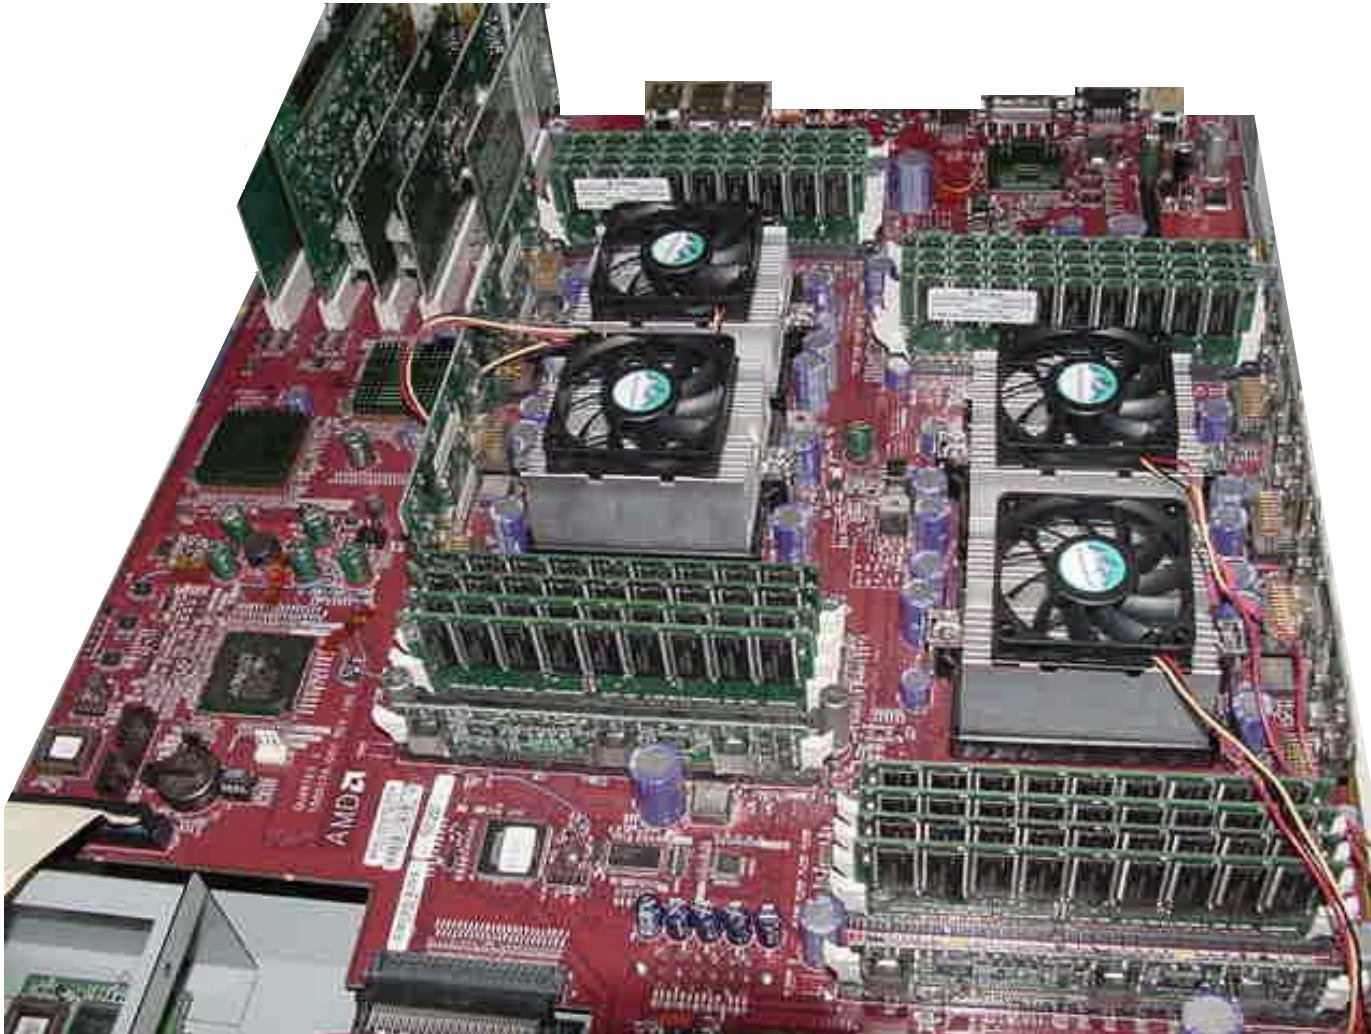
## "Glueless Multiprocessing"

- No chipset logic needed to connect processors
- HyperTransport™ technology links with ~6.4GBytes/sec bandwidth
- Memory BW and capacity and I/O capacity designed to grow with # CPUs

# AMD Opteron™ processor-based 4P Server



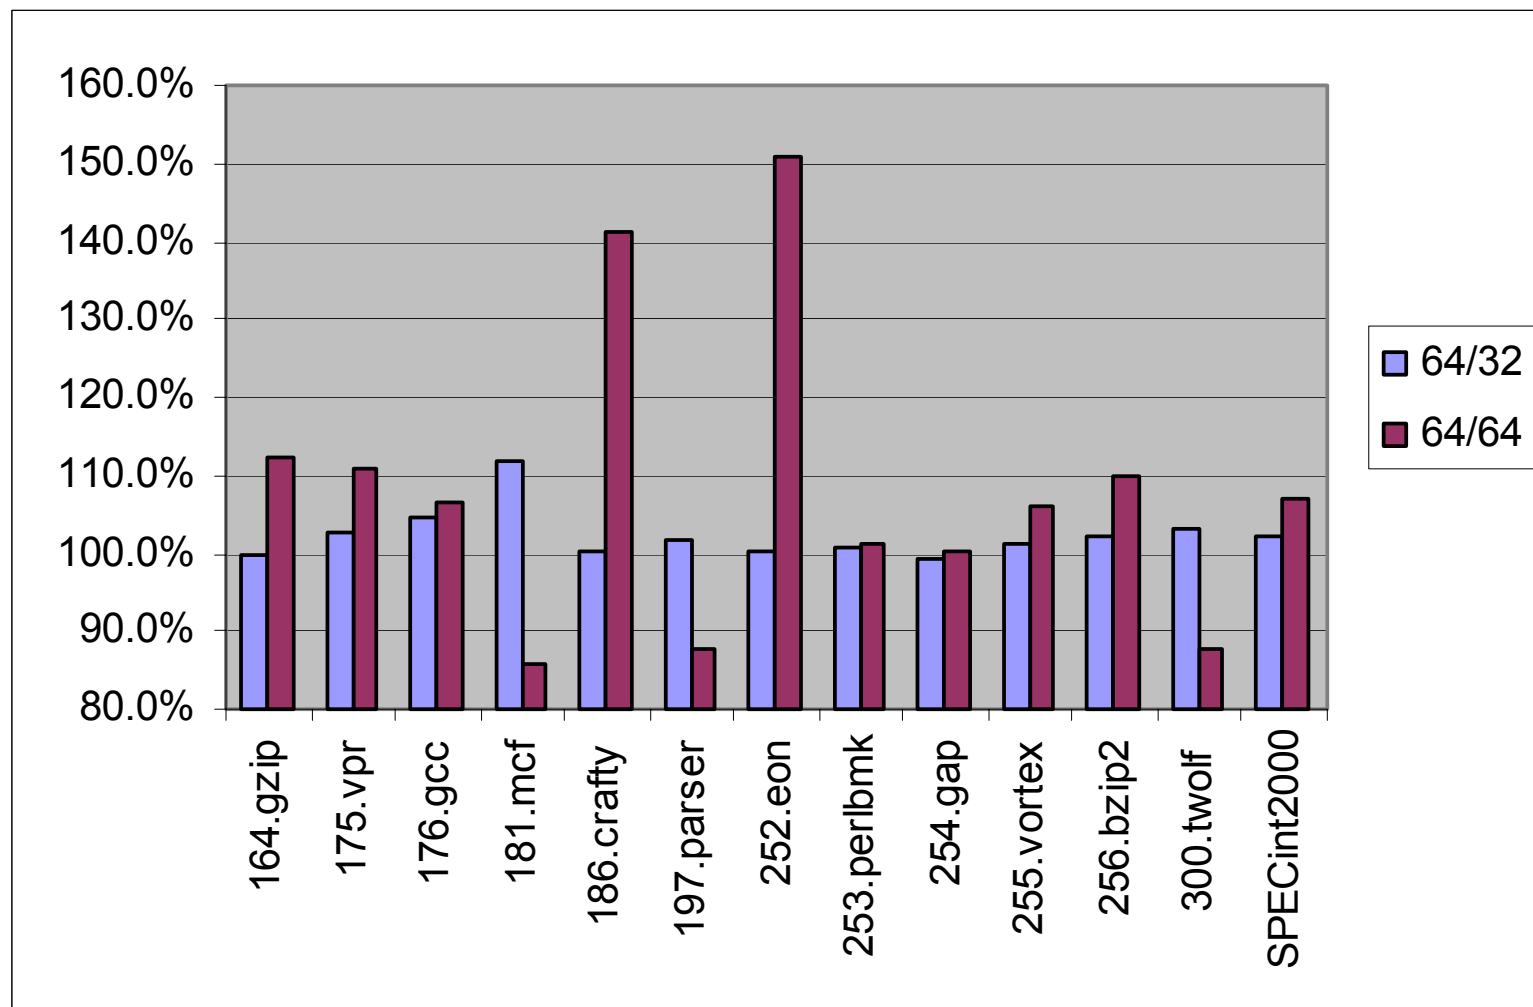
# 4P, 32GB AMD Opteron™ Processor System



# 4U, 4P AMD Opteron™ processor System



- GCC port alpha quality since Feb '01
  - Compiler generating alpha quality code in 50 man-months
  - Linux kernel ported in 60 man-months
  - Tool chain was straightforward port
- SpecInt2000 code quality, 64bits vs. 32 bits (using GCC 3.1.1)
  - average instruction length increased to 3.8 from 3.4 bytes
  - dynamic instruction count decreased by 10%
  - dynamic load count decreased by 26%
    - number of loads forwarded from recent stores substantially reduced
  - dynamic store count decreased by 36%
  - back to back register dependencies decreased by 10%





## The Portland Group Compiler Technology

- AMD and STMicroelectronics are working together to bring The Portland Group Compiler Technology to AMD64
  - Support will include
    - F90 & F77
      - Some F95 extensions also included
      - SPECcpu2000 explicitly supported
    - Optimized 32-bit and 64-bit code generation
    - Linux and Windows
    - OpenMP support
    - Full debugging support
  - STMicro will also be developing C and C++ compilers based on same code generation technology
  - Beta now, Production quality in 1H03

## • High Bandwidth

- 2P system is designed to achieve 7 GB/s aggregate memory Read bandwidth
- 4P system is designed to achieve 10 GB/s aggregate memory Read bandwidth
  - With data spread uniformly across the nodes

## • Low Latency

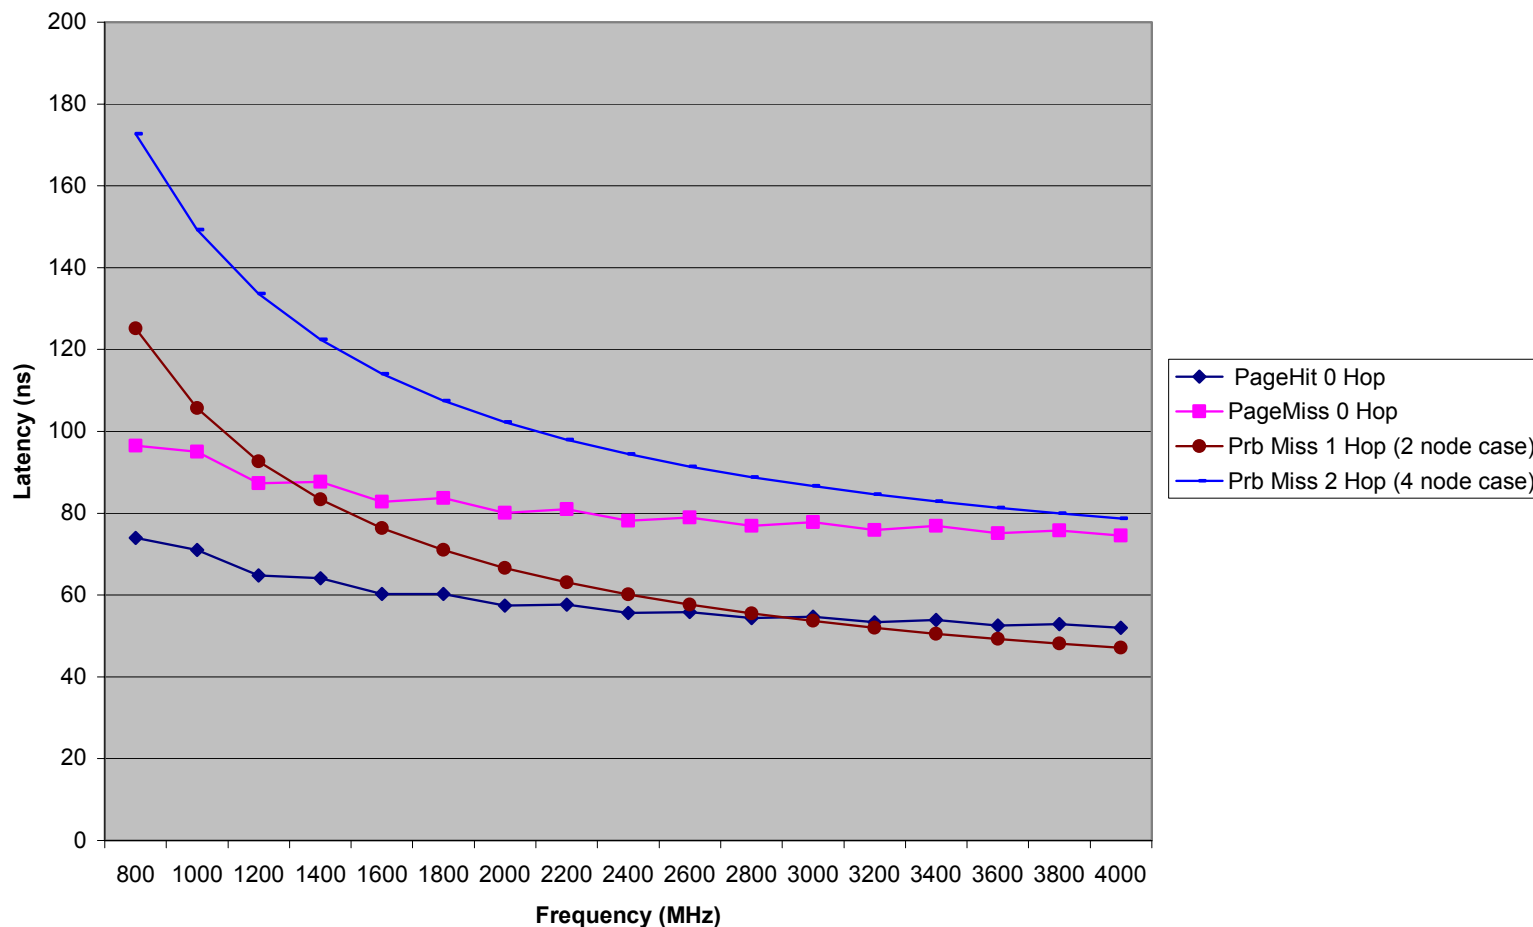
- Average 2P unloaded latency (page hit) is designed to be  $< 120$  ns
- Average 4P unloaded latency (page hit) is designed to be  $< 140$  ns
- Latency under load increases slowly due to excess Interconnect Bandwidth
- Latency shrinks quickly with increasing CPU clock speed and HyperTransport link speed

# Integrated Memory Controller

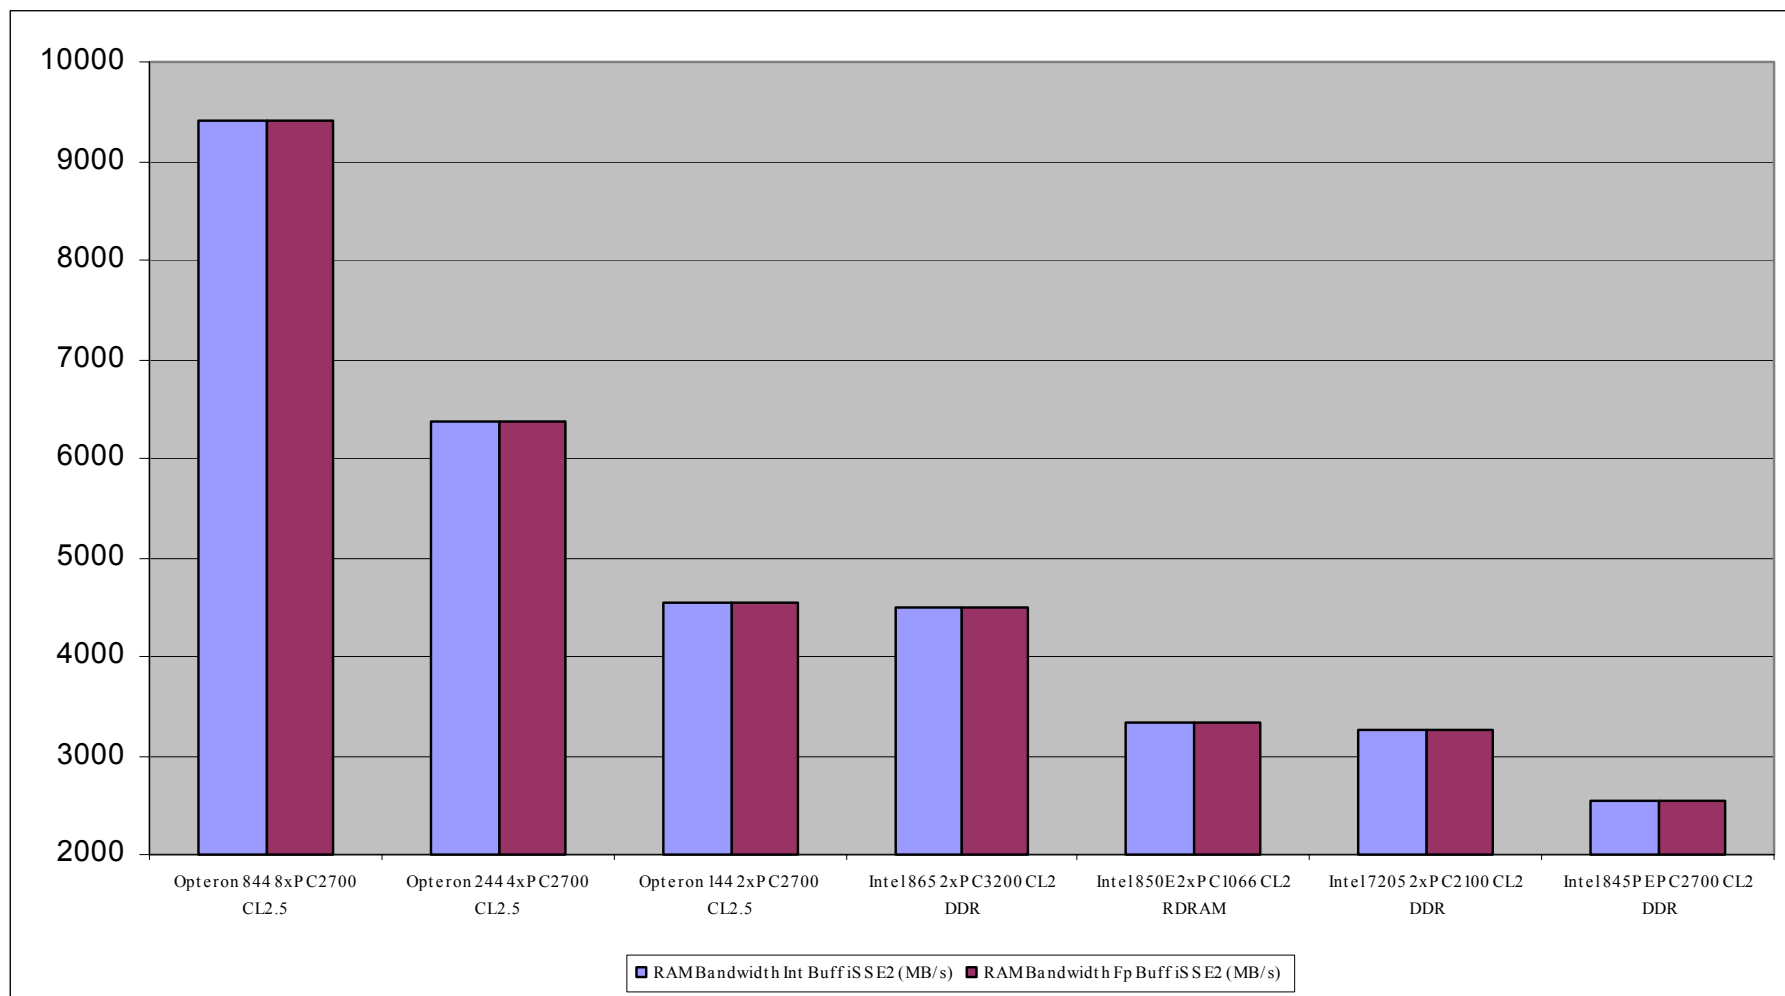
Latency (Local Memory Access, Registered Memory, CAS2)



Read Latency Accessing Local Memory, PC2100



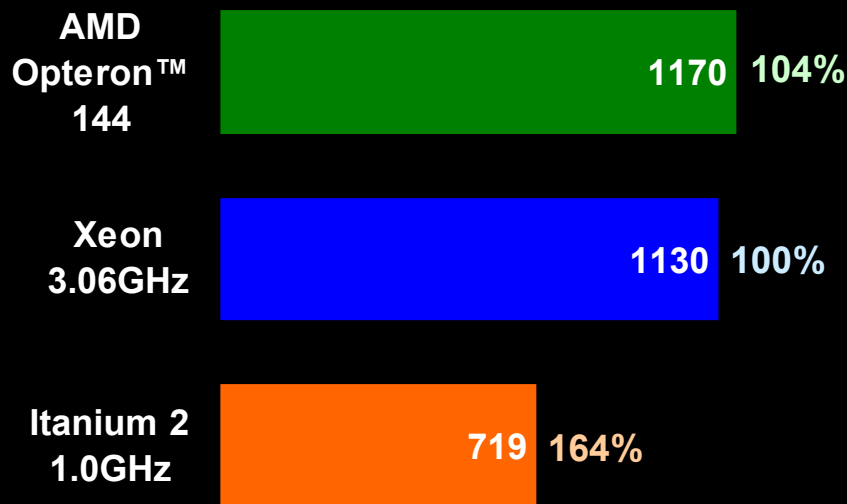
# Memory Bandwidth



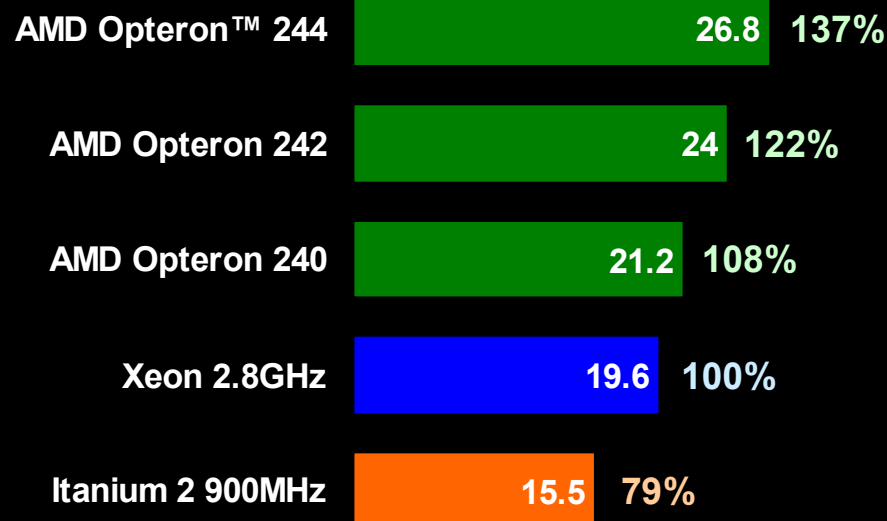
# Integer Performance



## SPECint®\_peak2000 Performance (Uniprocessor)



## SPECint®\_rate2000 Performance (Peak, 2P)



[www.amd.com/opteronperformance](http://www.amd.com/opteronperformance)

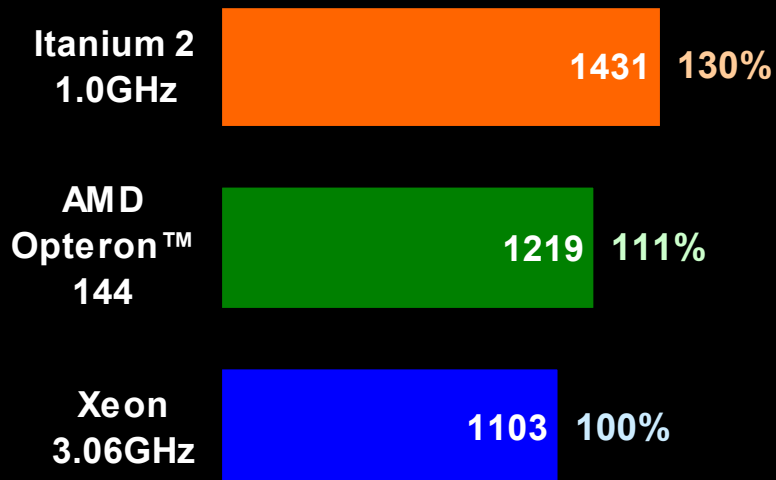


© 2003 Advanced Micro Devices. AMD, the AMD Arrow Logo, AMD Opteron and any combinations thereof are trademarks of Advanced Micro Devices. Microsoft and Windows are trademarks of Microsoft Corporation. SPEC and the benchmark named SPECint are registered trademarks of the Standard Performance Evaluation Corporation. Competitive benchmark results stated above reflect results published on [www.spec.org](http://www.spec.org) as of Apr 14, 2003. SPEC benchmark scores for AMD Opteron processor-based systems are under submission to the SPEC organization. For complete configuration information visit [www.spec.org](http://www.spec.org).

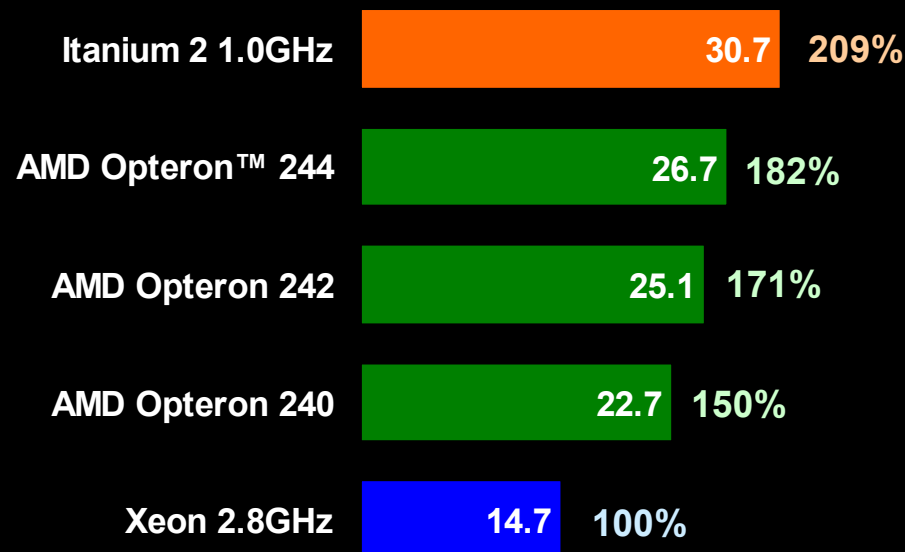
# Floating-Point Performance



## SPECfp®\_peak2000 Performance (Uniprocessor)



## SPECfp®\_rate2000 Performance (Peak, 2P)



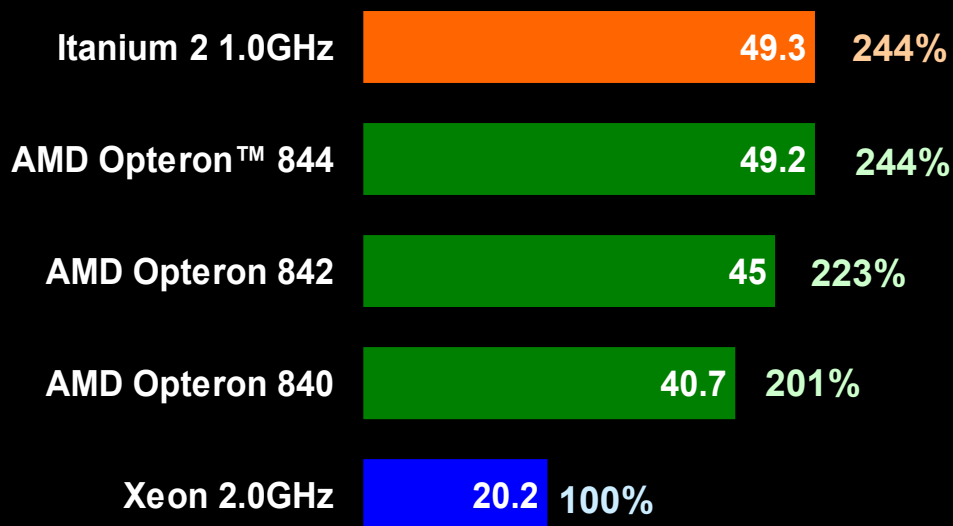
[www.amd.com/opteronperformance](http://www.amd.com/opteronperformance)



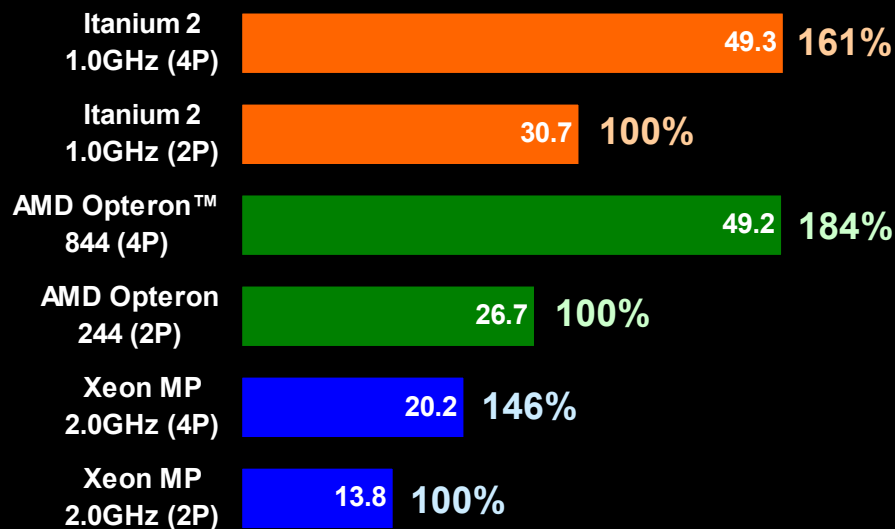
# Floating-Point Performance



## SPECfp®\_rate2000 Performance (Peak, 4P)



## SPECfp®\_rate2000 Performance and Scalability (Peak, 2-4P scaling)



[www.amd.com/opteronperformance](http://www.amd.com/opteronperformance)



# Web Server Performance



## SPECweb®99 Performance (2P Servers, Red Hat CA2)

AMD Opteron™ 244 6250 116%

AMD Opteron 242 5800 108%

Xeon  
3.06 GHz 5373 100%

AMD Opteron 240 5181 96%

Itanium 2 N/A

## SPECweb®99 Performance (4P Servers, Red Hat CA2)

AMD Opteron™ 844 10135 151%

AMD Opteron 842 9396 140%

AMD Opteron 840 8800 131%

Xeon MP  
2.0 GHz 6700 100%

Itanium 2 N/A

[www.amd.com/opteronperformance](http://www.amd.com/opteronperformance)

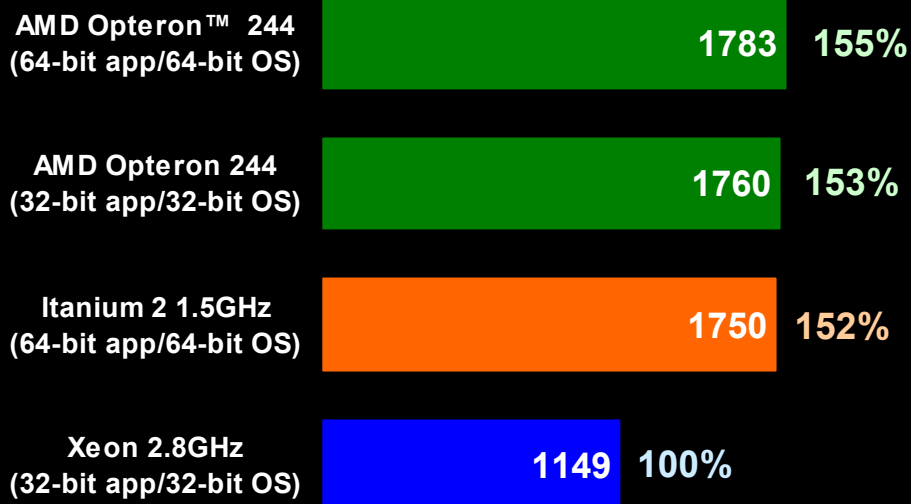


© 2003 Advanced Micro Devices. AMD, the AMD Arrow Logo, AMD Opteron and any combinations thereof are trademarks of Advanced Micro Devices. SPEC and the benchmark named SPECweb are registered trademarks of the Standard Performance Evaluation Corporation. Competitive benchmark results stated above reflect results published on [www.spec.org](http://www.spec.org) as of Apr 14, 2003. SPEC benchmark scores for AMD Opteron processor-based systems are under submission to the SPEC organization. For complete configuration information visit [www.spec.org](http://www.spec.org).

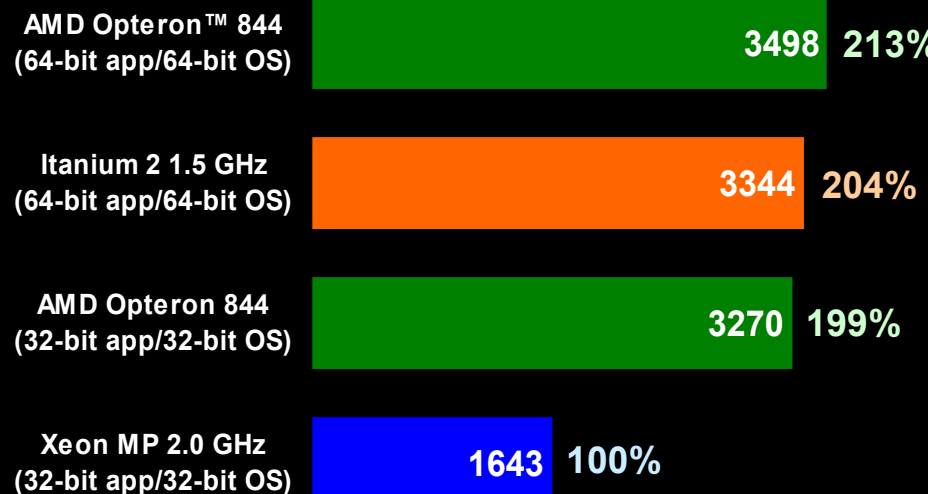
# Secure Web Server Performance



## SPECweb®99\_ssl Performance (2P Servers)



## SPECweb®99\_ssl Performance (4P Servers)



[www.amd.com/opteronperformance](http://www.amd.com/opteronperformance)



© 2003 Advanced Micro Devices. AMD, the AMD Arrow Logo, AMD Opteron and any combinations thereof are trademarks of Advanced Micro Devices. SPEC and the benchmark named SPECweb are registered trademarks of the Standard Performance Evaluation Corporation. Competitive benchmark results stated above reflect results published on [www.spec.org](http://www.spec.org) as of Apr 14, 2003. SPEC benchmark scores for AMD Opteron processor-based systems are under submission to the SPEC organization. For complete configuration information visit [www.spec.org](http://www.spec.org).

# Email Server Performance



## MMB2 Performance (2P Servers, Windows®)

AMD  
Opteron™  
244

11000 112%

Xeon  
2.8GHz

9800 100%

Itanium 2 **N/A**

## MMB2 Performance (4P Servers, Windows®)

AMD  
Opteron™  
844

15520 118%

Xeon MP  
2.0GHz

13200 100%

Itanium 2 **N/A**

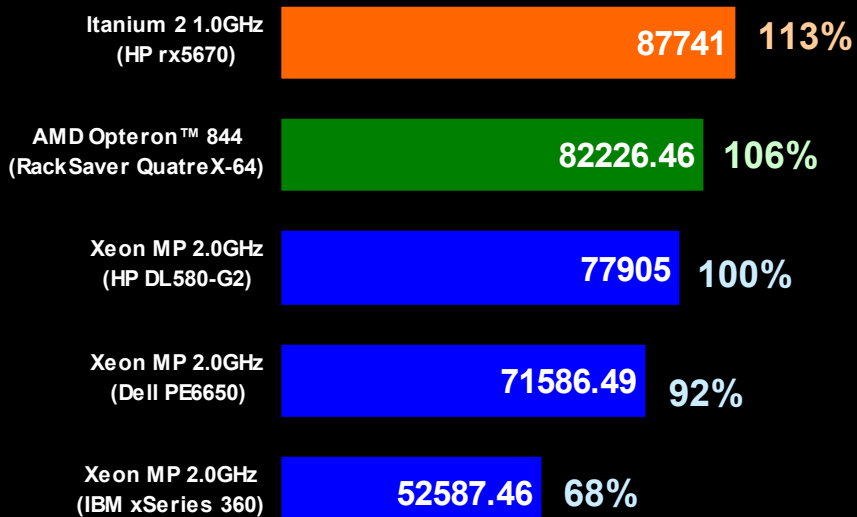
[www.amd.com/opteronperformance](http://www.amd.com/opteronperformance)



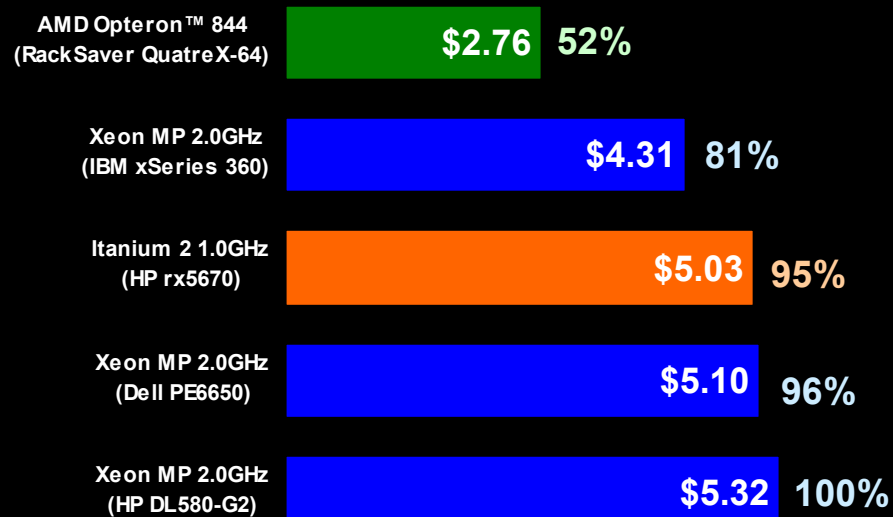
# Database Server Performance



## TPC-C Performance (4P Servers, tpmC, Windows®)



## TPC-C Price/Performance (4P Servers, \$/tpmC, Windows®)



[www.amd.com/opteronperformance](http://www.amd.com/opteronperformance)



# Linpack – Hot off the press



| AMD Opteron™ system                                   | # P | Rmax<br>(GFlops ) | Nmax<br>(order) | N1/2<br>(order) | Rpeak<br>(GFlops) | GFLO<br>Ps/Pr<br>oc | Rmax<br>Gflops /<br>Cycle | RPEAK/<br># Procs | Peak<br>Gflops /<br>Cycle | Rmax /<br>Peak |
|---|-----|-------------------|-----------------|-----------------|-------------------|---------------------|---------------------------|-------------------|---------------------------|----------------|
| 4P Melody Opteron 1.8GHz<br>2GB/proc PC2700 8GB Total | 4   | 11.99             | 28000           |                 | 14.4              | 3.00                | 1.665                     | 3.60              | 2.00                      | 83.3%          |
| 2P Melody Opteron 1.8GHz<br>2GB/proc PC2700 4GB Total | 2   | 6.009             | 19320           | 616             | 7.2               | 3.00                | 1.669                     | 3.60              | 2.00                      | 83.5%          |
| 1P Melody Opteron 1.8GHz 2GB<br>PC2700                | 1   | 3.042             | 14000           |                 | 3.6               | 3.04                | 1.690                     | 3.60              | 2.00                      | 84.5%          |

# CPU Design Clusters – From RISC to AMD64



- K6 was built entirely on Sparc, PA-RISC and Power machines
- K7/Athlon was built 50% on K6 running Linux
  - Few apps. Mostly only ran in house logic simulators
- K8/Opteron was built 80% on K7 running Linux
  - Many apps available. Only 64 bit apps conspicuously missing
- Hardware
  - Over 3000 Athlon CPUs doing back-end CAD work in California and Austin
  - Over 1500 Athlon CPUs doing front-end design world-wide
  - Non-AMD machines are used only for applications which require more memory than x86 is capable of addressing
- Software
  - Predominantly Linux based
  - Transitioning away from non-x86 based Unix (Solaris, HP-UX, etc.)
  - 64-bit software is run on non-AMD machines

- K9 will be taped out using **only** AMD Opteron Processors
- Hardware
  - Create a homogenous compute environment
  - Anticipate over 8000 AMD Opteron/Athlon CPUs doing back-end CAD work in Sunnyvale and Austin
  - Anticipate over 2000 AMD Opteron/Athlon CPUs doing front-end design world-wide
  - AMD **will not** use any non-AMD 32-bit or 64-bit hardware
- Software
  - 100% Linux/LSF based throughput cluster
  - 32-bit and 64-bit applications running side by side
  - Large memory applications will scale well on Opteron – **4P = 16-32 GB of RAM**

- The right instruction set
  - Excellent compatability
  - Excellent performance future
- The right system architecture
  - Great memory and IO capacity and bandwidth
  - Great memory latency
  - Simple “lego” system configuration
- A strong ecosystem of commodity HW and SW
  - Support chips, Software tools, motherboards
- Millions of 64 bit CPUs in 03
- 10s of millions of 64 bit CPUs in 04

- Allow more balanced scale-up/scale-out future
  - Remove 2P/4P cost barrier
  - And eventually 8P, 16P
- Re-create the workstation
  - Constrained by 32 bit x86 on one side and slow RISC processors w/o desktop software on the other
  - 2P, 16GB, 64 bit Workstation that runs Outlook, Powerpoint and Unreal Tournament
    - 64 bit portables in 04
- X86 forever (sorry ☺)

- Moore's law continues through the decade (and beyond)
  - 90nm, 65nm, 45nm, 30nm
  - 1 Billion transistors, 4 Billion transistors
  - Vertical integration
    - It will come, first for memory
    - Gigabyte on a die goes a long way to help memory wall
- Power is the biggest issue
  - Cache, Evaporators ☺
  - Metal gate, FinFet, Adiabatic clocks, etc
- CMP is good (and obvious)
- Threading is a mixed bag
  - Latency tolerance vs. Ahmdal's law and synchronization overhead
    - Long history
  - Certainly not for execution unit utilization

- Communication barrier
  - More fundamental than memory barrier
  - Even the speed of light doesn't help (much)
  - 3D helps a lot
- Single Chip Performance (a guess)

|        |               |                       |
|--------|---------------|-----------------------|
| – 2003 | 5 +/- 1 Gflop | Opteron, P4, iTanium2 |
| – 2005 | 12 Gflop      | 2 * 6GHz              |
| – 2006 | 24 Gflop      | 2P * 2 * 6GHz         |
| – 2007 | 36-72 Gflop   | 4 * 9GHz              |
| – 2008 | 144 Gflop     | 4P * 4 * 9GHz         |

# What Can You Do To Help



- Killer Apps that drive what you want
  - Games
  - Video compression/decompression
  - Face recognition as a ubiquitous app
- Keep the faith on COTS